
A PERFORMANCE ON-DEMAND APPROACH TO POWER-EFFICIENT COMPUTING

**David H. Albonesi
Sandhya Dwarkadas
Eby G. Friedman
Michael L. Scott**

**University of Rochester
Department of ECE; 411 Computer
Science's Building, Box 270231
Rochester, NY 14627-0231**

14 July 2004

Final Report

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION IS UNLIMITED.



**AIR FORCE RESEARCH LABORATORY
Space Vehicles Directorate
3550 Aberdeen Ave SE
AIR FORCE MATERIEL COMMAND
KIRTLAND AIR FORCE BASE, NM 87117-5776**

DTIC COPY

AFRL-VS-PS-TR-2004-1141

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data, does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report has been reviewed by the Public Affairs Office and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nationals.

If you change your address, wish to be removed from this mailing list, or your organization no longer employs the addressee, please notify AFRL/VSSE, 3550 Aberdeen Ave SE, Kirtland AFB, NM 87117-5776.

Do not return copies of this report unless contractual obligations or notice on a specific document requires its return.

This report has been approved for publication.

//signed//

JIM LYKE
Project Manager

//signed//

KIRT S. MOSER, DR-IV
Chief, Spacecraft Technology Division

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 14-07-2004		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 18-05-2000 to 14-07-2004	
4. TITLE AND SUBTITLE A Performance On-Demand Approach to Power-Efficient Computing				5a. CONTRACT NUMBER F29601-00-K-0182	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 62301E	
6. AUTHOR(S) David H. Albonesi, Sandhya Dwarkadas, Eby G. Friedman, Michael L. Scott				5d. PROJECT NUMBER DARP	
				5e. TASK NUMBER SE	
				5f. WORK UNIT NUMBER BL	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Rochester Dept of ECE; 411 Computer Science's Bldg, Box 270231 Rochester, NY 14627-0231				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory Space Vehicles Directorate 3550 Aberdeen Ave., SE Kirtland AFB, NM 87117-5776				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-VS-PS-TR-2004-1141	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Complexity-Adaptive Processing (CAP) addresses increasing microprocessor power dissipation through on-the-fly, low-cost hardware adaptation and related circuit techniques so as to better match hardware complexity and speed to application demands. Specific results include adaptive processing elements and hardware/software control techniques, a Multiple Clock Domain processor that saves energy via fine-grain voltage scaling, power-efficient issue queue and register file techniques, a low-leakage dynamic logic circuit and associated control logic for functional units, multi-threaded power and noise reduction, efficient on-chip dc-dc conversion and clock control circuits, low power domino logic and interface circuits, and interconnect width optimization for low power. Overall, a several-fold reduction in power is demonstrated via the collective application of these various techniques.					
15. SUBJECT TERMS Power-Aware Computing, Computer Architecture, Circuit Design					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Unlimited	18. NUMBER OF PAGES 20	19a. NAME OF RESPONSIBLE PERSON Dr. Jim Lyke
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code) (505) 846-5812

Contents

1	Project Objectives and Overview	1
2	Technical Accomplishments	1
2.1	Adaptive Techniques for Performance and Energy [9, 10, 13, 14, 15, 17, 18, 19, 21]	2
2.2	Multiple Clock Domain Microarchitecture [41, 42, 53, 55, 56]	3
2.3	Power-Efficient Issue Queues [14, 16]	5
2.4	High-Speed, Power-Aware Register Files [11, 12]	5
2.5	Reducing Static Power in Microprocessor Functional Units [22]	5
2.6	Multi-Threaded Processor Power and Noise Reduction [23, 24, 25]	6
2.7	Efficient On-Chip dc-dc Conversion [36, 37, 38, 39]	6
2.8	Low Power Domino Logic [32, 33, 34, 35]	7
2.9	On-Chip Power Distribution Network Optimization [43, 44, 45, 46, 47, 48]	7
2.10	Inductive Interconnect Width Optimization for Low Power [26, 27]	8
2.11	Low Power Voltage Interface Circuit [40]	8
2.12	Variable Clock Frequency Circuit [52]	8
2.13	Microarchitecture and Circuit Simulation Infrastructure	9
3	Ongoing Work	9
A	List of Patents	15

1 Project Objectives and Overview

The overall objective of the Complexity-Adaptive Processing (CAP) project is to provide on-the-fly, low-cost hardware adaptation so as to better match hardware complexity and speed to application demands, thereby dramatically improving power efficiency without unduly compromising performance. The CAP approach is to incorporate novel, low-intrusive feedback and control mechanisms into conventional microprocessors, so as to retain their high clock rates and high functional density while better matching their hardware resources to varying application phase characteristics. A combination of hardware and system software controls each element of performance and dynamic power: hardware complexity (switched capacitance), latency, clock frequency, and supply voltage. These elements are manifested as dynamic hardware structures and fine-grained clock frequency and voltage control circuits, and are controlled so as to meet performance objectives in the most power-efficient manner possible.

The dynamic hardware structures of the CAP project exploit the characteristics of major microprocessor hardware structures. In the very-deep-submicrometer regime, large on-chip RAM and CAM-based structures require repeaters in their global wires in order to minimize propagation delay. These repeaters are converted into low-overhead switches that electrically isolate individual sections of the structure, thereby allowing sections to be almost instantaneously turned on or off. The resulting dynamic hardware structures can be reorganized (e.g., resized) on-the-fly to match the different hardware requirements of different application phases.

As part of the CAP project, a Multiple Clock Domain (MCD) processor microarchitecture is investigated. In MCD, the processor is split into multiple domains, within which the frequency and supply voltage can be independently scaled. Synchronization circuits assure reliable communication among domains. In this manner, those domains that are not a performance bottleneck for a particular application phase can be run at lower frequency and voltage, thereby saving energy with tolerable performance impact. This fine-grain voltage scaling approach is effective across a wide range of general-purpose and embedded applications, in contrast to the limited utility of global voltage scaling.

Lower power organizational alternatives to complex hardware structures are also being developed as well as novel circuit and software techniques for power reduction. A defining characteristic of the CAP project is that it synergistically combines innovations at the circuits, architecture, and software levels.

Another defining characteristic of the CAP project is close ties with leading industry research laboratories in order to enable technology transfer. Indeed, two of the PhD students supported on this project are now researching power-aware microarchitecture and circuits as Research Staff Members in industry, one at the IBM T.J. Watson Research Center, and the other at Intel's Barcelona Research Center. (One other is interviewing at IBM Austin Research and T.J. Watson Research Labs, while three are Assistant Professors in academia, at the University of Wisconsin-Madison [starting this fall], the University of Utah, and Rochester Institute of Technology.) Both IBM and Intel provided their own funding for CAP research throughout most of the contract period, and continue to do so.

2 Technical Accomplishments

Under this contract, our team invented a wealth of novel approaches for reducing microprocessor power dissipation with minimal and area performance costs. These are briefly summarized below, and described in detail in the provided publications.

2.1 Adaptive Techniques for Performance and Energy [9, 10, 13, 14, 15, 17, 18, 19, 21]

Our group introduced what we termed *complexity adaptive processors* to the research community in 1998 at ISCA [1] and an associated workshop [2]. Since then, the field of *adaptive processing* [6] has been a very active area of research, both by our group [1, 2, 3, 4, 5, 8, 9, 10, 14, 15, 17, 18, 19, 21, 58, 59], as well as by many others [7, 20, 28, 29, 30, 31, 49, 50, 51, 57, 60, 61]. The adaptive techniques that our group developed under this contract are as follows:

- *Adaptive cache and TLB hierarchies* [9, 10]: A cache and TLB layout and design is devised that leverages repeater insertion to provide dynamic low-cost configurability, trading off size and speed on a per application phase basis. A novel configuration management algorithm is developed that dynamically detects phase changes and reacts to an application's hit and miss intolerance in order to improve memory hierarchy performance while taking energy consumption into consideration. When applied to a two-level cache and TLB hierarchy at $0.1\mu\text{m}$ technology, the result is an average 15% reduction in cycles per instruction (CPI), corresponding to an average 27% reduction in memory-CPI, across a broad class of applications compared to the best conventional two-level hierarchy of comparable size. Projecting to sub- $1\mu\text{m}$ technology design considerations that call for a three-level conventional cache hierarchy for performance reasons, we demonstrate that a configurable L2/L3 cache hierarchy coupled with a conventional L1 results in an average 43% reduction in memory hierarchy energy in addition to improved performance.
- *Adaptive issue queues* [14, 15, 18, 19]: We perform a circuit design of an issue queue for a superscalar processor that leverages transmission gate insertion to provide dynamic low-cost configurability of size and speed. A novel circuit structure dynamically gathers statistics of issue queue activity over intervals of instruction execution. These statistics are then used to change the size of an issue queue organization on-the-fly to improve issue queue energy and performance. When applied to a fixed, full-size issue queue structure, the result is up to a 70% reduction in energy dissipation. Using IBM process parameters and libraries used in a high-end processor, we determine that the complexity of the additional circuitry is almost negligible. Furthermore, self-timed techniques embedded in the adaptive scheme provide a 56% decrease in cycle time of the CAM array read of the issue queue when we change the adaptive issue queue size from 32 entries (largest possible) to 8 entries (smallest possible in our design).
- *Integrating multiple adaptive structures, including caches, issue queue, register files, and the Reorder Buffer* [21]: Prior adaptive hardware studies analyzed individual structures and their control. A common theme to these studies is exploration of the configuration space and use of system IPC as feedback to guide reconfiguration. However, when multiple structures adapt in concert, the number of possible configurations increases dramatically, and assigning causal effects to IPC change becomes problematic. To overcome this issue, we develop designs that are reconfigured solely on local behavior. We invent a novel cache design that permits direct calculation of efficient configurations. For buffer and queue structures, *limited histogramming* permits precise resizing control. When applying these techniques, we show energy savings of up to 70% on the individual structures, and savings averaging 30% overall for the portion of energy attributed to these structures, with only an average 2.1% performance cost.
- *Co-adaptive instruction fetch and issue* [17]: Front-end instruction delivery accounts for a significant fraction of the energy consumed in a dynamic superscalar processor. The issue queue in these processors serves two crucial roles: it bridges the front and back ends of the processor and serves as the window of instructions for the out-of-order engine. A mismatch between the front end producer rate and back end consumer rate, and between the supplied instruction window from the front end, and the

required instruction window to exploit the level of application parallelism, results in additional front-end energy, and increases the issue queue utilization. While the former increases overall processor energy consumption, the latter aggravates the issue queue hot spot problem.

We develop a complementary combination of fetch gating and issue queue adaptation to address both of these issues. We introduce an issue-centric fetch gating scheme based on issue queue utilization and application parallelism characteristics. Our scheme attempts to provide an instruction window size that matches the current parallelism characteristics of the application while maintaining enough queue entries to avoid back-end starvation. Compared to a conventional fetch gating scheme based on flow-rate matching, we demonstrate 20% better overall energy-delay with a 44% additional reduction in issue queue energy. We identify Icache energy savings as the largest contributor to the overall savings and quantify the sources of savings in this structure. We then couple this issue-driven fetch gating approach with an issue queue adaptation scheme based on queue utilization. While the fetch gating scheme provides a window of issue queue instructions appropriate to the level of program parallelism, the issue queue adaptation approach shuts down the remaining underutilized issue queue entries. Used in tandem, these complementary techniques yield a 20% greater issue queue energy savings than the addition of the savings from each technique applied in isolation. The result of this combined approach is a 6% overall energy-delay savings coupled with a 54% reduction in issue queue energy.

- *Energy-efficient adaptive clustered processors* [13]: Clustered microarchitectures are an attractive alternative to large monolithic superscalar designs due to their potential for higher clock rates in the face of increasingly wire-delay-constrained process technologies. As increasing transistor counts allow an increase in the number of clusters, thereby allowing more aggressive use of instruction-level parallelism (ILP), the inter-cluster communication increases as data values get spread across a wider area. As a result of the emergence of this trade-off between communication and parallelism, a subset of the total on-chip clusters is optimal for performance. To match the hardware to the application's needs, we use a robust algorithm to dynamically tune the clustered architecture. The algorithm, which is based on program metrics gathered at periodic intervals, achieves an 11% performance improvement on average over the best statically defined architecture. We also show that the use of additional hardware and reconfiguration at basic block boundaries can achieve average improvements of 15% while using on average four out of eight clusters, permitting these clusters to be turned off to save power when they are not needed. Our results demonstrate that reconfiguration provides an effective solution to the communication and parallelism trade-off inherent in the communication-bound processors of the future.

A subset of this work is summarized in our article in the IEEE Computer special issue on Power-Aware Computing [6].

2.2 Multiple Clock Domain Microarchitecture [41, 42, 53, 55, 56]

As clock frequency increases and feature size decreases, clock distribution and wire delays present a growing challenge to the designers of singly-clocked, globally synchronous systems. We develop an alternative approach, which we call a *Multiple Clock Domain (MCD)* processor, in which the chip is divided into several (coarse-grained) clock domains, within which independent voltage and frequency scaling can be performed [42, 56]. Boundaries between domains are chosen to exploit existing queues, thereby minimizing inter-domain synchronization costs. We propose four clock domains, corresponding to the front end (including L1 instruction cache), integer units, floating point units, and load-store units (including L1 data cache and L2 cache). We evaluate this design using a simulation infrastructure based on SimpleScalar and Wattch. In an attempt to quantify potential energy savings independent of any particular on-line control

strategy, we use off-line analysis of traces from a single-speed run of each of our benchmark applications to identify profitable reconfiguration points for a subsequent dynamic scaling run. Dynamic runs incorporate a detailed model of inter-domain synchronization delays, with latencies for intra-domain scaling similar to the whole-chip scaling latencies of Intel XScale and Transmeta LongRun technologies. Using applications from the MediaBench, Olden, and SPEC2000 benchmark suites, we obtain an average energy-delay product improvement of 20% with MCD compared to a modest 3% savings from voltage scaling a single clock and voltage system.

Subsequently, we invent an online, hardware-based, algorithm to dynamically control the frequency/voltage of MCD [53]. Our approach, which we call the *Attack/Decay Algorithm*, monitors differences in domain input queue utilization over intervals of operation. The algorithm adjusts the frequency and voltage of a domain if large differences are observed, and otherwise decays the frequency/voltage in small increments. Our algorithm achieves on average a 19.0% reduction in Energy Per Instruction (EPI), a 3.2% increase in Cycles Per Instruction (CPI), a 16.7% improvement in EnergyDelay Product, and a Power Savings to Performance Degradation ratio of 4.6. Traditional frequency/voltage scaling techniques which apply reductions globally to a fully synchronous processor achieve a Power Savings to Performance Degradation ratio of only 23. Our EnergyDelay Product improvement is 85.5% of that achieved using the prior offline algorithm [56].

We then devise techniques for automatic insertion of reconfiguration instructions into applications, using profile-driven binary rewriting [41]. Profile-based reconfiguration introduces the need for “training runs” prior to production use of a given application, but avoids the hardware complexity of on-line reconfiguration. It also has the potential to yield significantly greater energy savings. Experimental results (training on small data sets and then running on larger, alternative data sets) indicate that the profile-driven approach is more stable than hardware-based reconfiguration, and yields virtually all of the energy-delay improvement achieved via off-line analysis. Specifically, the approach yields an average 31% overall processor energy savings with only a 7% performance degradation, a result which compares very favorably with the near-ideal offline approach [56].

We also analyze a simulated Alpha 21264-like MCD microarchitecture in order to identify the architectural features of the processor that influence the less-than-expected performance degradation due to inter-domain synchronization [55]. We show that the out-of-order superscalar execution and decoupling features of a high performance microprocessor, which allow latency to be hidden, are the same features that reduce the performance degradation impact of the synchronization costs of an MCD processor. In the case of our Alpha 21264-like processor, up to 94% of the MCD synchronization delays are hidden and do not impact overall performance. In addition, we show that by adding out-of-order superscalar execution capabilities to a simpler microarchitecture, such as an Intel StrongARM-like processor, as much as 62% of the performance degradation caused by synchronization delays can be eliminated.

Finally, we combine our adaptive processing and MCD techniques in order to improve performance [54]. We explore “upsizing” hardware resources in order to improve performance relative to an aggressively clocked baseline processor. We use a variant of our MCD processor with four independently clocked domains. Each domain is streamlined with modest hardware structures for very high clock frequency. Key structures can then be upsized on demand to exploit more distant parallelism, improve branch prediction, or increase cache capacity. Although doing so requires decreasing the associated domain frequency, other domain frequencies are unaffected. Measuring across a broad suite of application benchmarks, we find that configuring just once per application increases performance by an average of 17.6% compared to the best fully synchronous design. When adapting to application phases, performance improves by over 20%.

A subset of this work is summarized in our article in the IEEE Micro special issue on the Top Picks from Microarchitecture Conferences [42]. Our ongoing MCD work is described in Section 3.

2.3 Power-Efficient Issue Queues [14, 16]

In addition to adaptive issue queues, we devise several other energy-efficient issue queue approaches. Several microprocessors, including the Alpha 21264 and POWER4, use a compacting latch-based issue queue design which has the advantage of simplicity of design and verification. The disadvantage of this structure, however, is its high power dissipation. We develop several different issue queue power optimization techniques that vary not only in their performance and power characteristics, but in how much they deviate from the baseline implementation. These techniques include fine-grain clock gating, non-compactation, a novel banking scheme, and dynamic adaptation. By developing and comparing techniques that build incrementally on the baseline design, as well as those that achieve higher power savings through a more significant redesign effort, we quantify the extra benefit the higher design cost techniques provide over their more straightforward counterparts.

2.4 High-Speed, Power-Aware Register Files [11, 12]

Modern superscalar processors use wide instruction issue widths and out-of-order execution in order to increase instruction-level parallelism (ILP). Because instructions must be committed in order so as to guarantee precise exceptions, increasing ILP implies increasing the sizes of structures such as the register file, issue queue, and reorder buffer. Simultaneously, cycle time constraints limit the sizes of these structures, resulting in conflicting design requirements.

To address these issues, we devise a novel microarchitecture designed to overcome the limitations of a register file size dictated by cycle time constraints [11]. Available registers are dynamically allocated between the *primary* program thread and a *future* thread. The *future* thread executes instructions when the *primary* thread is limited by resource availability. The *future* thread is not constrained by in-order commit requirements. It is therefore able to examine a much larger instruction window and jump far ahead to execute ready instructions. Results are communicated back to the *primary* thread by warming up the register file, instruction cache, data cache, and instruction reuse buffer, and by resolving branch mispredicts early. The proposed microarchitecture is able to get an overall speedup of 1.17 over the base processor for our benchmark set, with speedups of up to 1.64.

The number of physical registers within the processor has a direct impact on the size of the instruction window as most in-flight instructions require a new physical register at dispatch. A large multiported register file helps improve the ILP, but may have a detrimental effect on clock speed, especially in future wire-limited technologies. We propose a register file organization that reduces register file size and port requirements (thereby saving power) for a given amount of ILP [12]. We use a two-level register file organization to reduce register file size requirements, and a banked organization to reduce port requirements. We demonstrate empirically that the resulting register file organizations have reduced latency and (in the case of the banked organization) energy requirements for similar instructions per cycle (IPC) performance and improved instructions per second (IPS) performance in comparison to a conventional monolithic register file. These optimizations reduce register file power dissipation by more than a factor of four.

2.5 Reducing Static Power in Microprocessor Functional Units [22]

Static energy due to subthreshold leakage current is projected to become a major component of the total energy in high performance microprocessors. Many studies so far have examined and proposed techniques to reduce leakage in on-chip storage structures. In this study, static energy is reduced in the integer functional units by leveraging the unique qualities of dual threshold voltage domino logic.

Domino logic has desirable properties that greatly reduce leakage current while providing fast propagation times. However, due to the energy cost of entering the low leakage current state (*sleep mode*), domino

logic has thus far been used only for leakage reduction in the long-term standby mode. We examine the utility of the *sleep* mode (while considering the aforementioned costs) when idle times are relatively short, one to a few hundred cycles, as is often the case for functional units.

We develop an analytical energy model suitable for architecture-level analysis, and use the model to explore the interaction of the application and technology, and the effect on energy and performance as the underlying parameters are varied, on a set of benchmarks. Our results show that if the leakage approaches the magnitude as projected in the literature, even for short idle intervals as few as ten cycles, an aggressive policy of activating the sleep mode at every idle period performs well and a more complex control strategy may not be warranted. We devise a novel approach, called *Gradual Sleep*, to reduce the energy impact of using the *sleep* mode for smaller idle periods. The gradual sleep policy is able to optimally exploit the sleep mode state for various degrees of static power, permitting the same policy to be used as the design is scaled into more aggressive technology.

2.6 Multi-Threaded Processor Power and Noise Reduction [23, 24, 25]

The performance and power optimization of dynamic superscalar microprocessors requires striking a careful balance between exploiting parallelism and hardware simplification. Hardware structures which are needlessly complex may exacerbate critical timing paths and dissipate extra power. One such structure requiring careful design is the issue queue. In a Simultaneous Multi-Threading (SMT) processor, it is particularly challenging to achieve issue queue simplification due to the increased utilization of the queue afforded by multi-threading.

We invent new front-end policies that reduce the required integer and floating point issue queue sizes in SMT processors [25]. We explore both general policies as well as those directed towards alleviating a particular cause of issue queue inefficiency. Two policies are particularly effective and easily implementable. The first counts the number of instructions in the issue queue for each thread that were dispatched with one or more source operands unavailable. Instructions are not fetched for a given thread if its count is above a given threshold. The other policy predicts when a fetched load will miss in the data cache later in the pipeline, and maintains a count of such instructions for each thread. Again, no instruction fetching occurs for a thread whose count exceeds a threshold. For the same level of performance, the most effective combination of these policies reduces the issue queue occupancy by 33% for an SMT processor with appropriately-sized issue queue resources, resulting in a commensurate level of issue queue power savings without loss of performance.

SMT processors also exacerbate the inductive noise problem such that more expensive electronic solutions are required even with the use of previously proposed microarchitectural approaches. We use detailed microarchitectural simulation together with the Pentium 4 power delivery model to demonstrate the impact of SMT on inductive noise, and to identify thread-specific microarchitectural reasons for high noise occurrences [24]. We make the key observation that the presence of multiple threads actually provides an opportunity to mitigate the cyclical current fluctuations that cause noise, and propose the use of a prior performance enhancement technique to achieve this purpose. Specifically, we demonstrate that the judicious combination of flushing with damping dramatically improves performance and power efficiency for a given guaranteed noise limit.

Our ongoing work in power-efficient multi-threading is discussed in Section 3.

2.7 Efficient On-Chip dc-dc Conversion [36, 37, 38, 39]

A novel on-chip buck converter is designed and analyzed using Intel process parameters and circuit libraries. A high switching frequency is the key design parameter that simultaneously permits monolithic integration and high efficiency. A model of the parasitic impedances of a buck converter is developed. With this model,

a design space is determined that allows integration of active and passive devices on the same die for a target technology. An efficiency of 88.4% at a switching frequency of 477 MHz is demonstrated for a voltage conversion from 1.2-0.9 volts while supplying 9.5 A average current. The area occupied by the buck converter is 12.6 mm² assuming an 80-nm CMOS technology. An estimate of the efficiency is shown to be within 2.4% of simulation at the target design point. Full integration of a high-efficiency buck converter on the same die with a dual- V_{DD} microprocessor is demonstrated to be feasible.

A low-voltage-swing MOSFET gate drive technique is proposed for enhancing the efficiency characteristics of the high-frequency-switching dc-dc converter. The parasitic power dissipation of a dc-dc converter is reduced by lowering the voltage swing of the power transistor gate drivers. A comprehensive circuit model of the parasitic impedances of a monolithic buck converter is presented. Closed-form expressions for the total power dissipation of a low-swing buck converter are proposed. The effect of reducing the MOSFET gate voltage swings is explored with the proposed circuit model. A range of design parameters is evaluated, permitting the development of a design space for full integration of active and passive devices of a low-swing buck converter on the same die, for a target CMOS technology. The optimum gate voltage swing of a power MOSFET that maximizes efficiency is lower than a standard full voltage swing. An efficiency of 88% at a switching frequency of 102 MHz is achieved for a voltage conversion from 1.8 to 0.9 V with a low-swing dc-dc converter based on a 0.18 μ m CMOS technology. The power dissipation of a low-swing dc-dc converter is reduced by 27.9% as compared to a standard full-swing dc-dc converter.

2.8 Low Power Domino Logic [32, 33, 34, 35]

A circuit technique is presented for reducing the subthreshold leakage energy consumption of domino logic circuits. Sleep switch transistors are proposed to place an idle dual threshold voltage domino logic circuit into a low leakage state. The circuit technique enhances the effectiveness of a dual threshold voltage CMOS technology to reduce the subthreshold leakage current by strongly turning off all of the high threshold voltage transistors. The sleep switch circuit technique significantly reduces the subthreshold leakage energy as compared to both standard low-threshold voltage and dual threshold voltage domino logic circuits. A domino adder enters and leaves a low leakage sleep mode within a single clock cycle. The energy overhead of the circuit technique is low, justifying the activation of the proposed sleep scheme by providing a net savings in total power consumption during short idle periods.

Furthermore, a variable threshold voltage keeper circuit technique is proposed for simultaneous power reduction and speed enhancement of domino logic circuits. The threshold voltage of a keeper transistor is dynamically modified during circuit operation to reduce contention current without sacrificing noise immunity. The variable threshold voltage keeper circuit technique enhances circuit evaluation speed by up to 60% while reducing power dissipation by 35% as compared to a standard domino (SD) logic circuit. The keeper size can be increased with the proposed technique while preserving the same delay or power characteristics as compared to a SD circuit. The proposed domino logic circuit technique offers 14% higher noise immunity as compared to a SD circuit with the same evaluation delay characteristics. Forward body biasing the keeper transistor is also proposed for improved noise immunity as compared to a SD circuit with the same keeper size. It is shown that by applying forward and reverse body biased keeper circuit techniques, the noise immunity and evaluation speed of domino logic circuits are simultaneously enhanced.

2.9 On-Chip Power Distribution Network Optimization [43, 44, 45, 46, 47, 48]

The design of power distribution networks in high-performance integrated circuits has become significantly more challenging with recent advances in process technologies. The inductive characteristics of several types of gridded power distribution networks are described. The inductance extraction program FastHenry is used to evaluate the inductive properties of grid structured interconnect. In power distribution grids with al-

ternating power and ground lines, the inductance is shown to vary linearly with grid length and inversely linearly with the number of lines in the grid. The inductance is also relatively constant with frequency in these grid structures. These properties permit the efficient estimation of the inductive characteristics of power distribution grids. To optimize the process of allocating on-chip metal resources, inductance/area/resistance tradeoffs in high speed performance distribution grids are explored. Two tradeoff scenarios in power grids with alternating power and ground lines are considered.

Furthermore, as on-chip currents exceed tens of amperes and circuit clock periods are reduced well below a nanosecond, the signal integrity of on-chip power supply has become a primary concern in the integrated circuit design. The scaling behavior of the inductive and resistance voltage drops across the on-chip power distribution networks is analyzed. The existing work on power distribution noise scaling is reviewed and extended to include the scaling behavior of the inductance of the on-chip global power distribution networks in high-performance flip-chip packaged integrated circuits. As the dimensions of the on-chip devices are scaled by S , where $S > 1$, the resistive voltage drop across the power grids remains constant and the inductive voltage drop increases by S , if the metal thickness is maintained constant. Consequently, the signal-to-noise ratio decreases by S in the case of resistive noise and by S^2 in the case of inductive noise. As compared to the constant metal thickness scenario, ideal interconnect scaling of the global power grid mitigates the unfavorable scaling of the inductive noise but exacerbates the scaling of resistive noise by a factor of S . On-chip inductive noise will, therefore, become of greater significance with technology scaling. Careful tradeoffs between the resistance and inductance of the power distribution networks will be necessary in nanometer technologies to achieve minimum power supply noise.

2.10 Inductive Interconnect Width Optimization for Low Power [26, 27]

The width of an interconnect line affects the total power consumed by a circuit. A tradeoff exists, however, between the dynamic power and the short-circuit power in determining the width of inductive interconnects. The optimum line width that minimizes the total transient power dissipation is determined. A closed form solution for the optimum width with an error of less than 6% is presented. For a specific set of line parameters and resistivities, the power is reduced by almost 80% as compared to a minimum wire width. Considering the driver size in the design process, the optimum wire and driver size that minimizes the total transient power is also determined. Furthermore, the use of similar optimization techniques in repeated lines results in a 65% reduction in power and 97% reduction in delay.

2.11 Low Power Voltage Interface Circuit [40]

A bi-directional CMOS voltage interface circuit is proposed for applications that require signal transfer between two circuits operating at different voltage levels. The circuit can also be used as a level converter at the driver and receiver ends of long interconnect lines for low swing applications. The operation of the voltage interface circuit is verified by both simulation and experimental test circuits. The proposed voltage interface circuit operates at high speed while offering significant power savings of up to 95% as compared to existing schemes.

2.12 Variable Clock Frequency Circuit [52]

A circuit to dynamically reconfigure the clock frequency of a synchronous digital system according to the changing needs of the application is developed. The circuit changes the clock frequency with a minimal time penalty and offers glitch free, reliable operation.

2.13 Microarchitecture and Circuit Simulation Infrastructure

In order to perform our research, we developed a sophisticated architecture, circuit, and application profiling infrastructure for performance, power, and noise modeling. This toolset is readily available to the research community and has been adopted by several groups.

3 Ongoing Work

The following summarizes our ongoing work in adaptive processing and the development of realizable MCD processors:

- *Simultaneous performance and energy optimization.* Our efforts in these two areas have been largely focused on reducing energy with minimal performance loss. Our work in [54] demonstrated that performance can be improved using adaptive processing within an MCD processor. However, this effort exclusively focused on performance and ignored energy efficiency; we also used adaptivity in only a subset of the potential candidates. Moreover, we have not begun to explore how to combine this approach with fine-grain dynamic frequency scaling [41, 53] within MCD. We believe that exploiting the natural synergy between these approaches will yield both increased performance and greater energy efficiency.
- *Adaptive processing and multi-threading.* Thus far, we have explored single-threaded architectures, whereas the use of SMT is rapidly gaining momentum in the server, desktop, and embedded marketplaces. We believe that the use of multiple threads makes for an even more compelling case for adaptive processing and MCD, due to the variation in the number of threads that may be running at any given time. For instance, many parallel applications have a substantial sequential component, and this single thread may run sub-optimally on an SMT supporting four threads due to the particular trade off made between resource size and clock speed at design time. Alternatively, multi-threaded performance may be traded off at design time by the need to have good single thread performance. Adaptive processing within an MCD design can be used to optimize the microarchitecture to the number of active threads and their characteristics.
- *Optimizing both dynamic and leakage energy.* Although there have been isolated efforts to reduce leakage energy using adaptation, our adaptive processing and MCD efforts to date have largely addressed dynamic power. A comprehensive study of how best to optimize combined dynamic and leakage energy needs to be undertaken. As part of this effort, efficient dynamic voltage scaling dc-dc converters, adaptive body bias generators, and voltage interface circuits are being developed.
- *Reducing the number of required processor cores.* The ubiquity of computing is leading to the development of many customized processor designs geared towards a particular class of applications. This requires either many processor design teams producing several custom-designed cores in parallel (thereby greatly increasing design costs) or the development of synthesized cores that are performance and energy sub-optimal. An open question is whether a single, custom designed, adaptive MCD processor, constrained so as not to unduly increase design time and die area, can provide a compelling technical and economic alternative to several custom core designs. We intend to perform a comparative evaluation of what environments and to what degree an adaptive MCD processor might provide an advantage over designing multiple custom cores.
- *Design simplifications.* Issues of design and verification complexity may easily override impressive performance and energy savings. We have identified design simplifications that still provide good results, in order to ease the practical implementation of these ideas. We have already made good headway in this effort with our IBM collaborators, who have shown great interest in the MCD design.

In addition, we are continuing our work in energy-efficient multi-threaded processors. We have developed a simulation infrastructure for the exploration of Clustered Multi-Threaded (CMT) processors, in which multiple threads share the resources of a clustered microarchitecture in a dynamic manner. Our preliminary results are very promising, showing a significant reduction in power compared to SMT while yielding competitive performance, and were recently presented at Intel and IBM.

References

- [1] D.H. Albonesi. Dynamic IPC/clock rate optimization. *Proceedings of the 25th International Symposium on Computer Architecture*, pages 282–292, June 1998.
- [2] D.H. Albonesi. The inherent energy efficiency of complexity-adaptive processors. *Proceedings of the 1998 Power-Driven Microarchitecture Workshop*, pages 107–112, June 1998.
- [3] D.H. Albonesi. An architectural and circuit-level approach to improving the energy efficiency of microprocessor memory structures. *Proceedings of the 10th International Conference on VLSI*, pages 192–205, December 1999.
- [4] D.H. Albonesi. Selective cache ways: On-demand cache resource allocation. *Proceedings of the 32nd International Symposium on Microarchitecture*, pages 248–259, November 1999.
- [5] D.H. Albonesi. Selective cache ways: On-demand cache resource allocation. *Journal of Instruction-Level Parallelism*, Vol. 2, 2000.
- [6] D.H. Albonesi et al. Dynamically tuning processor resources with adaptive processing. *IEEE Computer*, 36(12):49–58, December 2003.
- [7] R. Iris Bahar and Srilatha Manne. Power and Energy Reduction Via Pipeline Balancing. In *Proceedings of ISCA-28*, July 2001.
- [8] R. Balasubramonian, D.H. Albonesi, A. Buyuktosunoglu, and S. Dwarkadas. Dynamic memory hierarchy performance optimization. *Workshop on Solving the Memory Wall Problem*, June 2000.
- [9] R. Balasubramonian, D.H. Albonesi, A. Buyuktosunoglu, and S. Dwarkadas. Memory hierarchy re-configuration for energy and performance in general-purpose processor architectures. *Proceedings of the 33rd International Symposium on Microarchitecture*, pages 245–257, December 2000.
- [10] R. Balasubramonian, D.H. Albonesi, A. Buyuktosunoglu, and S. Dwarkadas. A dynamically tunable memory hierarchy. *IEEE Transactions on Computers*, 52(10):243–258, October 2003.
- [11] R. Balasubramonian, S. Dwarkadas, and D.H. Albonesi. Dynamically allocating processor resources between nearby and distant ilp. *Proceedings of the 28th International Symposium on Computer Architecture*, pages 26–37, June 2001.
- [12] R. Balasubramonian, S. Dwarkadas, and D.H. Albonesi. Reducing the complexity of the register file in dynamic superscalar processors. *Proceedings of the 34th International Symposium on Microarchitecture*, pages 237–248, December 2001.
- [13] R. Balasubramonian, S. Dwarkadas, and D.H. Albonesi. Dynamically managing the communication-parallelism trade-off in future clustered processors. *Proceedings of the 30th International Symposium on Computer Architecture*, pages 275–286, June 2003.
- [14] A. Buyuktosunoglu, D.H. Albonesi, P. Bose, P. Cook, and S. Schuster. Tradeoffs in power-efficient issue queue design. *International Symposium on Low Power Electronics and Design*, pages 184–189, August 2002.
- [15] A. Buyuktosunoglu, D.H. Albonesi, S. Schuster, D. Brooks, P. Bose, and P. Cook. Power-efficient issue queue design. *Power Aware Computing*, pages 37–60, 2002.
- [16] A. Buyuktosunoglu, A. El-Moursy, and D.H. Albonesi. An oldest-first selection logic implementation for non-compacting issue queues. *Proceedings of the IEEE International ASIC/SOC Conference*, pages 31–35, September 2002.

- [17] A. Buyuktosunoglu, T. Karkhanis, D.H. Albonesi, and P. Bose. Energy efficient co-adaptive instruction fetch and issue. *Proceedings of the 30th International Symposium on Computer Architecture*, pages 147–156, June 2003.
- [18] A. Buyuktosunoglu, S. Schuster, D. Brooks, P. Bose, P. Cook, and D.H. Albonesi. An adaptive issue queue for reduced power at high performance. *Workshop on Power-Aware Computer Systems*, November 2000.
- [19] A. Buyuktosunoglu, S. Schuster, D. Brooks, P. Bose, P. Cook, and D.H. Albonesi. A circuit level implementation of an adaptive issue queue for power-aware microprocessors. *Proceedings of the 11th Great Lakes Symposium on VLSI*, March 2001.
- [20] A.S. Dhodapkar and J.E. Smith. Managing multi-configuration hardware via dynamic working set analysis. *International Symposium on Computer Architecture*, May 2002.
- [21] S. Dropsho, A. Buyuktosunoglu, R. Balasubramonian, D.H. Albonesi, S. Dwarkadas, G. Semeraro, G. Magklis, and M.L. Scott. Integrating adaptive on-chip storage structures for reduced dynamic power. *Proceedings of the 11th International Conference on Parallel Architectures and Compilation Techniques*, pages 141–152, September 2002.
- [22] S. Dropsho, V. Kursun, D.H. Albonesi, S. Dwarkadas, and E.G. Friedman. Managing static leakage energy in microprocessor functional units. *Proceedings of the 35th International Symposium on Microarchitecture*, pages 321–332, November 2002.
- [23] W. El-Essawy and D.H. Albonesi. A microarchitectural-level step-power analysis tool. *International Symposium on Low Power Electronics and Design*, pages 263–266, August 2002.
- [24] W. El-Essawy and D.H. Albonesi. Mitigating inductive noise in smt processors. *International Symposium on Low Power Electronics and Design*, August 2004.
- [25] A. El-Moursy and D.H. Albonesi. Front-end policies for improved issue efficiency in smt processors. *Proceedings of the 9th International Symposium on High Performance Computer Architecture*, pages 31–40, February 2003.
- [26] M.A. El-Moursy and E.G. Friedman. Optimizing inductive interconnect for low power. *Proceedings of the International Workshop on System-on-Chip for Real-Time Applications*, July 2002.
- [27] M.A. El-Moursy and E.G. Friedman. Optimizing wire sizing of RLC interconnect with repeaters. *Proceedings of the Great Lakes Symposium on VLSI*, pages 27–32, April 2003.
- [28] D. Folegnani and A. Gonzalez. Energy-effective issue logic. *International Symposium on Computer Architecture*, June 2001.
- [29] M. Huang, J. Reneau, S.M. Yoo, and J. Torrellas. A Framework for Dynamic Energy Efficiency and Temperature Management. In *Proceedings of MICRO-33*, pages 202–213, Dec 2000.
- [30] M.C. Huang et al. Positional adaptation of processors: application to energy reduction. *International Symposium on Computer Architecture*, June 2003.
- [31] Christopher J. Hughes, Jayanth Srinivasan, and Sarita V.Adve. Saving energy with architectural and frequency adaptations for multimedia applications. *Proceedings of the 34th International Symposium on Microarchitecture*, December 2001.
- [32] V. Kursun, , and E.G. Friedman. Variable threshold voltage keeper for contention reduction in dynamic circuits. *Proceedings of the IEEE International ASIC/SOC Conference*, pages 314–318, September 2002.

- [33] V. Kursun, , and E.G. Friedman. Domino logic with variable threshold voltage keeper. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 11(6):1080–1093, December 2003.
- [34] V. Kursun and E.G. Friedman. Low swing dual threshold voltage domino logic. *Proceedings of the ACM/SIGDA Great Lakes Symposium on VLSI*, April 2002.
- [35] V. Kursun and E.G. Friedman. Sleep switch dual threshold voltage domino logic with reduced standby leakage current. *IEEE Transactions on Very Large Scale Integration*, 12(5):485–496, May 2004.
- [36] V. Kursun, S.G. Narendra, V.K. De, and E.G. Friedman. Efficiency analysis of a high frequency buck converter for on-chip integration with a dual-VDD microprocessor. *Proceedings of the European Solid-State Circuits Conference*, September 2002.
- [37] V. Kursun, S.G. Narendra, V.K. De, and E.G. Friedman. Analysis of buck converters for on-chip integration with a dual supply voltage microprocessor. *IEEE Transactions on Very Large Scale Integration*, 11(3):514–522, June 2003.
- [38] V. Kursun, S.G. Narendra, V.K. De, and E.G. Friedman. Monolithic dc-dc converter analysis and MOS-FET gate voltage optimization. *IEEE/ACM International Symposium on Quality Electronic Design*, March 2003.
- [39] V. Kursun, S.G. Narendra, V.K. De, and E.G. Friedman. Low-voltage-swing monolithic dc-dc conversion. *IEEE Transactions on Circuits and Systems–II*, 51(5):241–248, May 2004.
- [40] V. Kursun, R.M. Secareanu, and E.G. Friedman. CMOS voltage interface circuit for low power systems. *Proceedings of the IEEE International Conference on Electronics, Circuits, and Systems*, May 2002.
- [41] G. Magklis, M.L. Scott, G. Semeraro, D.H. Albonesi, and S. Dropsho. Profile-based dynamic voltage and frequency scaling for a multiple clock domain microprocessor. *Proceedings of the 30th International Symposium on Computer Architecture*, pages 14–25, June 2003.
- [42] G. Magklis, G. Semeraro, D.H. Albonesi, S.G. Dropsho, S. Dwarkadas, , and M.L. Scott. Dynamic frequency and voltage scaling for a multiple-clock-domain microprocessor. *IEEE Computer*, 23(6):62–68, November/December 2003.
- [43] A. Mezhiba and E.G. Friedman. Inductance/area/resistance tradeoffs in high performance power distribution grids. *Proceedings of the IEEE International Symposium on Circuits and Systems*, pages 1.101–1.104, May 2002.
- [44] A. Mezhiba and E.G. Friedman. Inductive properties of high-performance power distribution grids. *IEEE Transactions on Very Large Scale Integration*, 10(6):762–776, December 2002.
- [45] A. Mezhiba and E.G. Friedman. Properties of on-chip inductive current loops. *Proceedings of the Great Lakes Symposium on VLSI*, pages 12–17, April 2002.
- [46] A. Mezhiba and E.G. Friedman. Scaling trends of on-chip power distribution noise. *Proceedings of the IEEE International Workshop on System-Level Interconnect Prediction*, pages 47–53, April 2002.
- [47] A. Mezhiba and E.G. Friedman. Variation of inductance with frequency in high performance power distribution grids. *Proceedings of the IEEE International ASIC/SOC Conference*, pages 421–425, September 2002.
- [48] A. Mezhiba and E.G. Friedman. Scaling trends of on-chip power distribution noise. *IEEE Transactions on Very Large Scale Integration*, 12(4):386–394, April 2004.

- [49] D. Ponomarev et al. Dynamic allocation of datapath resources for low power. *International Symposium on Microarchitecture*, December 2001.
- [50] M.D. Powell et al. An energy-efficient high-performance deep-submicron instruction cache. *IEEE Transactions on VLSI Systems*, February 2001.
- [51] R. Sasanka, C.J. Hughes, and S.V. Adve. Joint local and global hardware adaptations for energy. *International Conference on Architecture Support for Programming Languages and Operating Systems*, October 2002.
- [52] R.M. Secareanu, D.H. Albonesi, and E.G. Friedman. A dynamic reconfigurable clock generator. *Proceedings of the IEEE ASIC Conference*, pages 330–333, September 2001.
- [53] G. Semeraro, D.H. Albonesi, S.G. Dropsho, G. Magklis, S. Dwarkadas, and M.L. Scott. Dynamic frequency and voltage control for a multiple clock domain microarchitecture. *Proceedings of the 35th International Symposium on Microarchitecture*, pages 356–367, November 2002.
- [54] G. Semeraro, D.H. Albonesi, S.G. Dropsho, G. Magklis, and M.L. Scott. Improving application performance by dynamically balancing speed and complexity in a GALS microarchitecture. *Workshop on Application Specific Processors*, December 2003.
- [55] G. Semeraro, D.H. Albonesi, G. Magklis, M.L. Scott, S.G. Dropsho, and S. Dwarkadas. Hiding synchronization delays in a GALS processor microarchitecture. *Proceedings of the 10th International Symposium on Asynchronous Circuits and Systems*, April 2004.
- [56] G. Semeraro, G. Magklis, R. Balasubramonian, D.H. Albonesi, S. Dwarkadas, and M.L. Scott. Dynamic frequency and voltage control for a multiple clock domain microarchitecture. *Proceedings of the 8th International Symposium on High Performance Computer Architecture*, pages 29–40, February 2002.
- [57] T. Sherwood, A. Sair, and B. Calder. Phase tracking and prediction. *Proceedings of the 30th International Symposium on Computer Architecture*, pages 336–347, June 2003.
- [58] B. Xu and D.H. Albonesi. A methodology for the analysis of dynamic application parallelism and its application to reconfigurable computing. *Proceedings of the SPIE International Symposium on Reconfigurable Technology: FPGAs for Computing and Applications*, pages 78–86, September 1999.
- [59] B. Xu and D.H. Albonesi. Runtime reconfiguration techniques for efficient general purpose computation. *IEEE Design and Test of Computers (Special Issue on Configurable Computing Machines)*, pages 42–52, January 2000.
- [60] Se-Hyun Yang, Michael D. Powell, Babak Falsafi, and T.N. Vijaykumar. Exploiting choice in resizable cache design to optimize deep-submicron processor energy-delay. *Proceedings of the International Symposium on High-Performance Computer Architecture*, February 2002.
- [61] C. Zhang, F. Vahid, and W. Najjar. A highly configurable cache architecture for embedded systems. *International Symposium on Computer Architecture*, June 2003.

A List of Patents

- Dynamically Reconfigurable Memory Hierarchy, number 6,684,298, R. Balasubramonian, D.H. Albonesi, A. Buyuktosunoglu, and S. Dwarkadas, issued 1/27/04.
- Multiple Clock Domain Microprocessor, D.H. Albonesi, G. Semeraro, G. Magklis, M. Scott, R. Balasubramonian, and S. Dwarkadas, filed, 1/23/04.
- Method of Adaptive Control and Resizing of Issue Queue Structures Within an Out-of-order Superscalar Processor for Reduced Power at High Performance, A. Buyuktosunoglu, S. Schuster, D. Brooks, P. Bose, P. Cook, and D.H. Albonesi, filed, 10/15/01.
- Dynamic Data Dependence Tracking, L. Chen, S. Dropsho, and D.H. Albonesi, provisional filed, 2004.
- Dynamic Management of the Communication-Parallelism Trade-off in Clustered Microarchitectures, R. Balasubramonian, S. Dwarkadas, and D.H. Albonesi, provisional filed, 2003.

DISTRIBUTION LIST

DTIC/OCP 8725 John J. Kingman Rd, Suite 0944 Ft Belvoir, VA 22060-6218	1 cy
AFRL/VSIL Kirtland AFB, NM 87117-5776	1 cy
AFRL/VSIH Kirtland AFB, NM 87117-5776	1 cy
University of Rochester Department of ECE 411 Computer Science's Bldg Box 270231 Rochester, NY 14627-0231	1 cy
Official Record Copy AFRL/VSSE, Dr. Jim Lyke	1 cy